



PCT/EP2004/052502

16 NOV 2004



The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

REC'D 13 JAN 2005
WIPO PCT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

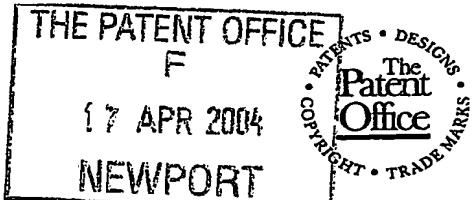
Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 20 October 2004

John
PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1 (a) OR (b)

BEST AVAILABLE COPY



19APR04 E889280-1 D02982
P01/7700 0.00-0408570.0 ACCOUNT CHA
The Patent Office

Request for grant of a patent

(See the notes on the back of this form. You can also get
an explanatory leaflet from the Patent Office to help you fill in
this form)

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference

P746GB

2. Patent application number

(The Patent Office will fill this part in)

17 APR 2004

0408570.0

3. Full name, address and postcode of the or of each applicant (underline all surnames)

UNIVERSITY OF YORK
HESLINGTON
YORK
YO1 5DD
GB

Patents ADP number (if you know it)

4169746001

If the applicant is a corporate body, give the
country/state of its incorporation

United Kingdom

4. Title of the invention

IMAGE RECOGNITION

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom
to which all correspondence should be sent
(including the postcode)

STANLEYS

Kings Court
12 King Street
LEEDS
LS1 2HL

Patents ADP number (if you know it)

07902190001

6. Priority: Complete this section if you are declaring priority from one or more earlier patent applications, filed in the last 12 months.

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. Divisionals, etc: Complete this section only if this application is a divisional application or resulted from an entitlement dispute (see note f)

Number of earlier UK application
(day / month / year)

8. Is a Patents Form 7/77 (Statement of inventorship and of right to grant of a patent) required in support of this request?

YES

Answer YES if:

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an
applicant, or
- c) any named applicant is a corporate body.

Otherwise answer NO (See note d)

Patents Form 1/77

9. Accompanying documents: A patent application must include a description of the invention. Not counting duplicates, please enter the number of pages of each item accompanying this form:

Continuation sheets of this form

Description	14
Claim(s)	<i>8</i>
Abstract	
Drawing(s)	

10. If you are also filing any of the following, state how many against each item.

Priority documents
Translations of priority documents
Statement of inventorship and right to grant of a patent (Patents Form 7/77)
Request for a preliminary examination and search (Patents Form 9/77)
Request for a substantive examination (Patents Form 10/77)

Any other documents (please specify)

Copy of GB 03 23662.7

11. I/We request the grant of a patent on the basis of this application.

STANLEYS

15 April 2004

Signature(s)

Stanley

Date

12. Name, daytime telephone number and e-mail address, if any, of person to contact in the United Kingdom

David Stanley

01481 824411

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered YES in part 8, a Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- Part 7 should only be completed when a divisional application is being made under section 15(4), or when an application is being made under section 8(3), 12(6) or 37(4) following an entitlement dispute. By completing part 7 you are requesting that this application takes the same filing date as an earlier UK application. If you want the new application to have the same priority date(s) as the earlier UK application, you should also complete part 6 with the priority details.

IMAGE RECOGNITION

This invention relates to the recognition of images, and is concerned, particularly although not exclusively, with the recognition of natural images.

By "natural image" is meant an image of an object that occurs naturally – for example, an optical image such as a photograph, as well as images of other wavelengths – such as x-ray and infra-red, by way of example. The natural image may be recorded and/or subsequently processed by digital means, but is in contrast to an image – or image data – that is generated or synthesised by computer or other artificial means.

The recognition of natural images can be desirable for many reasons. For example, distinctive landscapes and buildings can be recognised, to assist in the identification of geographical locations. The recognition of human faces can be useful for identification and security purposes. The recognition of valuable animals such as racehorses may be very useful for identification purposes.

In this specification, we present in preferred embodiments of the invention a new approach to face recognition.

Preferred embodiments of the present invention may be combined with techniques disclosed in our pending application GB0323662.7, a copy of the specification and drawings of which is attached, and to which the reader's attention is directed.

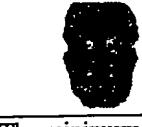
Previous work [1,2,3,4,5,6,7,8] has shown that the use of 3D face models is able to overcome some of the problems associated with 2D face recognition. Firstly, by relying on geometric shape, rather than colour and texture information, systems become invariant to lighting conditions. Secondly, the ability to rotate a facial structure in three-dimensional space, allowing for compensation of variations in pose, aids those methods requiring alignment prior to recognition.

Finally, the additional discriminatory information available in the facial surface structure, not available from two-dimensional images, provides additional cues for recognition.

It has also been shown that the use of pre-processing techniques applied 5 prior to training and recognition, in which distinguishing features are made more explicit, environmental effects are normalised and noise content is reduced, can significantly improve recognition accuracy [1, 9, 10]. However, the focus of previous research has been on identifying the optimum surface representation, with little regard for the advantages offered by each individual surface 10 representation. We suggest that different surface representations may be specifically suited to different capture conditions or certain facial characteristics, despite having a general weakness for overall recognition. For example, curvature representations may aid recognition by making the system more robust to inaccuracies in 3D orientation, yet be highly sensitive to noise. Another 15 representation may enhance nose shape, but loose the relative positions of facial features. The benefit of using multiple eigenspaces has previously been examined by Pentland et al [11], in which specialist eigenspaces were constructed for various facial orientations and local facial regions, from which cumulative match scores were able to reduce error rates. Our approach differs in that we extract and 20 combine individual dimensions, creating a single unified surface space. This approach has been shown to work effectively when applied to two-dimensional images by Heseltine et al [12].

Here we analyse and evaluate a range of 3D face recognition systems, each utilising a different surface representation of the facial structure, in an 25 attempt to identify and isolate the advantages offered by each representation. Focusing on the fishersurface method of face recognition, we propose a means of selecting and extracting components from the surface subspace produced by each system, such that they may be combined into a unified surface space.

Prior to training and testing, 3D face models are converted into one of the following surface representations. This is done by firstly orientating the 3D face model to face directly forwards, then projecting into a depth map. The surfaces in the table below are then derived by pre-processing of depth maps.

Horizontal Derivative	Vertical Derivative	Horizontal Derivative 2	Vertical Derivative 2
	-1 1		-1 1
Applies the 2x1 kernel to compute the horizontal derivative	Applies the 1x2 kernel to compute the vertical derivative	Horizontal gradient over a greater horizontal distance	Vertical gradient over a greater vertical distance
Laplacian	Sobel X	Sobel Y	Sobel Magnitude
	0 1 0 1 -4 1 0 1 0		-1 0 1 -2 0 2 -1 0 1
An isotropic measure of the second spatial derivative	Application of the horizontal sobel derivative filter	Application of the vertical sobel derivative filter	The magnitude of Sobel X and Y combined.
Horizontal Curvature	Vertical Curvature	Curvature Magnitude	Curve Type
			
Applies sobel X twice to calculate the second horizontal derivative	Applies sobel Y twice to calculate the second vertical derivative	The magnitude of the vertical and horizontal curvatures	Segmentation of the surface into 8 discreet curvature types
Min Curvature	Max Curvature	Abs Min Curvature	Abs Max Curvature
			
The minimum of the horizontal and vertical curvature values	The maximum of the horizontal and vertical curvature values	The minimum of the absolute horizontal and vertical curvatures	The maximum of the absolute horizontal and vertical curvatures

We give here a brief explanation of the fisherface method of face recognition, as described by Belhumeur et al [13] and how it is applied to three-

dimensional face surfaces, termed the fishersurface method. We apply both principal component analysis and linear discriminant analysis to surface representations of 3D face models, producing a subspace projection matrix, similar to that used in the eigenface [11] and eigensurface [1] methods. However, 5 the fishersurface method is able to take advantage of 'within-class' information, minimising variation between multiple face models of the same person, yet still maximising class separation. To accomplish this, we expand the training set to contain multiple examples of each subject, describing the variance of a person's face structure (due to influences such as facial expression and head orientation), 10 from one face model to another, as shown in equation 1.

$$\text{Training Set} = \underbrace{\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5, \Gamma_6, \Gamma_7, \Gamma_8, \Gamma_9, \Gamma_{10}, \Gamma_{11}, \Gamma_{12}, \Gamma_{13}, \dots \Gamma_M\}}_{X_1} \quad (1)$$
$$\underbrace{\phantom{\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5, \Gamma_6, \Gamma_7, \Gamma_8, \Gamma_9, \Gamma_{10}, \Gamma_{11}, \Gamma_{12}, \Gamma_{13}, \dots \Gamma_M\}}}_{X_2} \quad \underbrace{\phantom{\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5, \Gamma_6, \Gamma_7, \Gamma_8, \Gamma_9, \Gamma_{10}, \Gamma_{11}, \Gamma_{12}, \Gamma_{13}, \dots \Gamma_M\}}}_{X_3} \quad \underbrace{\phantom{\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5, \Gamma_6, \Gamma_7, \Gamma_8, \Gamma_9, \Gamma_{10}, \Gamma_{11}, \Gamma_{12}, \Gamma_{13}, \dots \Gamma_M\}}}_{X_4} \quad \underbrace{\phantom{\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5, \Gamma_6, \Gamma_7, \Gamma_8, \Gamma_9, \Gamma_{10}, \Gamma_{11}, \Gamma_{12}, \Gamma_{13}, \dots \Gamma_M\}}}_{X_c}$$

15 Where Γ_i is a facial surface and the training set is partitioned into c classes, such that each surface in each class X_i is of the same person and no single person is present in more than one class. We continue by computing three scatter matrices, representing the within-class (S_w), between-class (S_B) and total (S_T) distribution of the training set throughout surface space, shown in equation 2.

$$S_T = \sum_{i=1}^M (\Gamma_i - \Psi)(\Gamma_i - \Psi)^T \quad S_B = \sum_{i=1}^c |X_i| (\Psi_i - \Psi)(\Psi_i - \Psi)^T \quad S_w = \sum_{i=1}^c \sum_{\Gamma \in X_i} (\Gamma_i - \Psi_i)(\Gamma_i - \Psi_i)^T \quad (2)$$

Where $\Psi = \frac{1}{M} \sum_{i=1}^M \Gamma_i$ is the average surface of the entire training set, and $\Psi_i = \frac{1}{|X_i|} \sum_{\Gamma \in X_i} \Gamma_i$, the average of class X_i . By performing PCA using the total scatter matrix S_T , and taking the top $M-c$ principal components, we produce a projection matrix U_{PCA} , used to reduce dimensionality of the within-class scatter matrix, ensuring it is non-singular before computing the top $c-1$ (in this case 49) eigenvectors of the reduced scatter matrix ratio, U_{FID} as shown in equation 3.

$$U_{fd} = \arg \max_U \left(\frac{|U^T U_{pca}^T S_B U_{pca} U|}{|U^T U_{pca}^T S_W U_{pca} U|} \right). \quad (3)$$

Finally, the matrix U_f is calculated as shown in equation 4, such that it may project a face surface into a reduced surface space of $c-1$ dimensions, in which the between-class scatter is maximised for all c classes, while the within-class scatter is minimised for each class X_i .

$$U_f = U_{fd} U_{pca} \quad (4)$$

Once the matrix U_f has been constructed it is used in much the same way as the projection matrix in the eigenface and eigensurface systems, reducing dimensionality of face surface vectors from 5330 to just 49 ($c-1$) elements. Again, like the eigenface system, the components of the projection matrix can be viewed as images, as shown in figure X for the depth map surface space.

Once surface space has been defined, we project a facial surface into surface space by a simple matrix multiplication using the matrix U_f , as shown in equation 5.

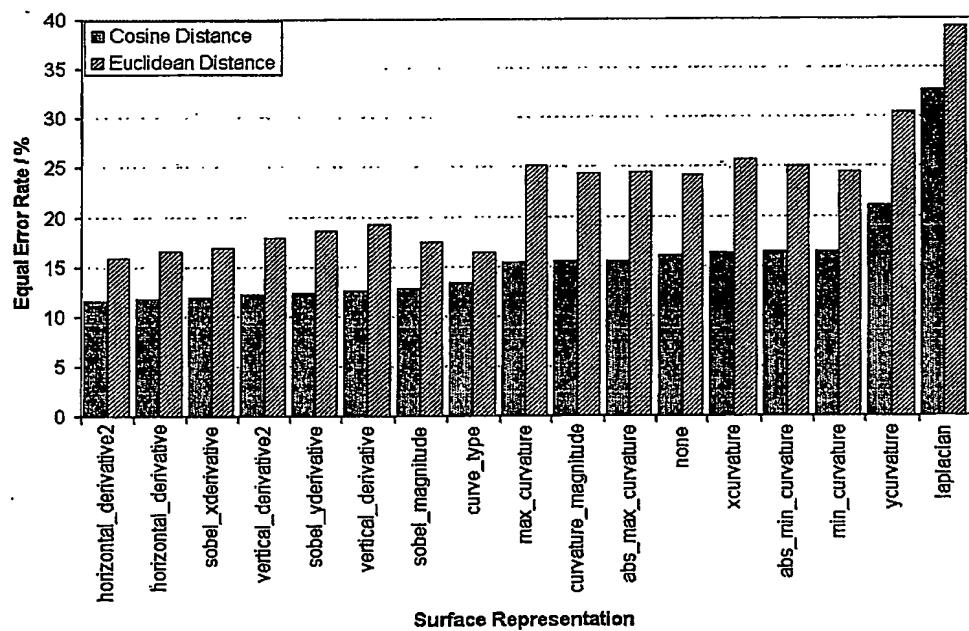
$$\omega_k = u_k^T (\Gamma - \Psi) \quad \text{for } k = 1 \dots c-1. \quad (5)$$

where u_k is the k th eigenvector and ω_k is the k th weight in the vector $\Omega^T = [\omega_1, \omega_2, \omega_3, \dots \omega_M]$. The $c-1$ coefficients represent the contribution of each respective fishersurface to the original facial surface structure. The vector Ω is taken as the 'face-key' representing a person's facial structure in reduced dimensionality surface space and compared using either euclidean or cosine distance metrics as shown in equation 6.

$$d_{euclidean} = \|\Omega_a - \Omega_b\| \quad d_{cosine} = 1 - \frac{\Omega_a^T \Omega_b}{\|\Omega_a\| \|\Omega_b\|} \quad (6)$$

5 An acceptance (the two facial surfaces match) or rejection (the two surfaces do not match) is determined by applying a threshold to the distance calculated. Any comparison producing a distance value below the threshold is considered an acceptance.

10 Here we analyse the surface spaces produced when various facial surface representations are used with the fishersurface method. We begin by providing results showing the range of error rates produced when using various surface representations. The figure below clearly shows that the choice of surface representation has a significant effect on the effectiveness of the fishersurface method, with horizontal gradient representations providing the lowest equal error rates (EER, the error when FAR equals FRR).



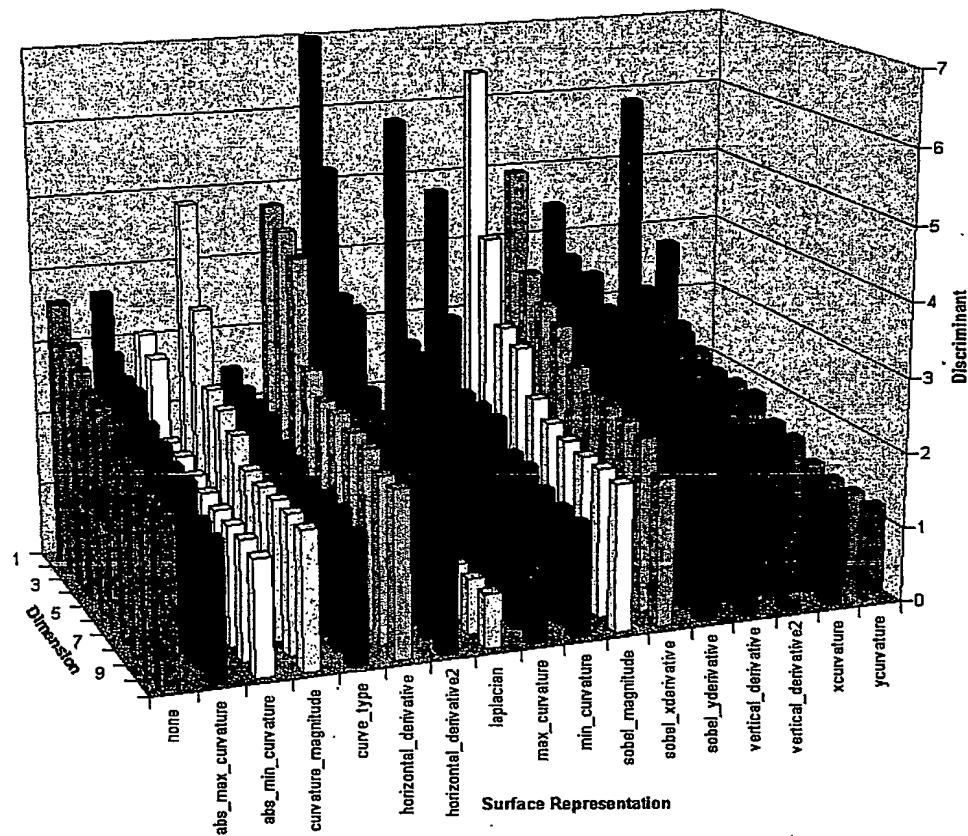
15 However, the superiority of the horizontal gradient representations does not suggest that the vertical gradient and curvature representations are of no use whatsoever and although the discriminatory information provided by these

representations may not be as robust and distinguishing, that is not to say they wouldn't make a positive contribution to the information already available in the horizontal gradient representations. We now carry out further investigation into the discriminating ability of each surface space by applying Fisher's Linear
5 Discriminant (FLD), as used by Gordon [3] to analyse 3D face features, to individual components (single dimensions) of each surface space. Focusing on a single face space dimension we calculate the discriminant d , describing the discriminating power of that dimension, between c people.

$$d = \frac{\sum_{i=1}^c (m_i - m)^2}{\sum_{i=1}^c \frac{1}{|\Phi_i|} \sum_{x \in \Phi_i} (x - m_i)^2}$$

10

Where m is the mean value of that dimension in the face-keys, m_i the within-class mean of class i and Φ_i the set of vector elements taken from the face-keys of class i . Applying the above equation to the assortment of surface space systems generated using each facial surface representation, we see a wide range of
15 discriminant values describing the distinguishing ability of each individual dimension, as shown below for the top ten most discriminating dimensions for each surface representation.



It is clear that although some surface representations do not perform well in the face recognition tests, producing high EERs (for example `min_curvature`), some of their face-key components do contain highly discriminatory information. We hypothesise that the reason for these highly discriminating anomalies, in an otherwise ineffective subspace, is that a certain surface representation may be particularly suited to a single discriminating factor, such as nose shape or jaw structure, but is not effective when used as a more general classifier. Therefore, if we were able to isolate these few useful qualities from the more specialised subspaces, they could be used to make a positive contribution to a generally more effective surface space, reducing error rates further.

Here we describe how the analysis methods discussed in above are used to combine multiple face recognition systems. Firstly, we need to address the problem of prioritising surface space dimensions. Because the average magnitude and deviation of face-key vectors from a range of systems are likely to differ by

some orders of magnitude, certain dimensions will have a greater influence than others, even if the discriminating abilities are evenly matched. To compensate for this effect, we normalise moments by dividing each face-key element by its within-class standard deviation. However, in normalising these dimensions we have also 5 removed any prioritisation, such that all face space components are considered equal. Although not a problem when applied to a single surface space, when combining multiple dimensions we would ideally wish to give greater precedence to the more reliable components. Otherwise the situation is likely to arise when a large number of less discriminating (but still useful) dimensions begin to outweigh 10 the fewer more discriminating ones, diminishing their influence on the verification operation and hence increasing error rates. We showed how FLD could be used to measure the discriminating ability of a single dimension from any given face space. We now apply this discriminant value d as a weighting for each face space dimension, prioritising those dimensions with the highest discriminating ability.

15 With this weighting scheme applied to all face-keys produced by each system, we can begin to combine dimensions into a single unified surface space. In order to combine multiple dimensions from a range of surface spaces, we require some criterion to decide which dimensions to combine. It is not enough to rely purely on the discriminant value itself, as this only gives us an indication of 20 the discriminating ability of that dimension alone, without any indication of whether the inclusion of this dimension would benefit the existing set of dimensions. If an existing surface space already provides a certain amount of discriminatory ability, it would be of little benefit (or could even be detrimental) if we were to introduce an additional dimension describing a feature already present 25 within the existing set.

Previous investigations [12] have used FLD, applied to a combined eigenspace in order to predict its effectiveness when used for recognition. Additional dimensions are then introduced if they result in a greater discriminant value. Such a method has been shown to produce an 2D eigenspace combination

able to achieve significantly lower error rates in 2D face recognition, although Heseltine et al also note that using the EER would likely provide better results, although processing time would be extremely long. However, with a more efficient combination algorithm we now take that approach, such that the criterion required 5 for a new dimension to be introduced to an existing surface space is a resultant increase in the EER.

Combined surface space = first dimension of current optimum system

Calculate *EER* of *combined surface space*

For each surface space system:

For each *dimension* of surface space:

Concatenate new *dimension* onto *combined surface space*

Calculate *EER* of *combined surface space*

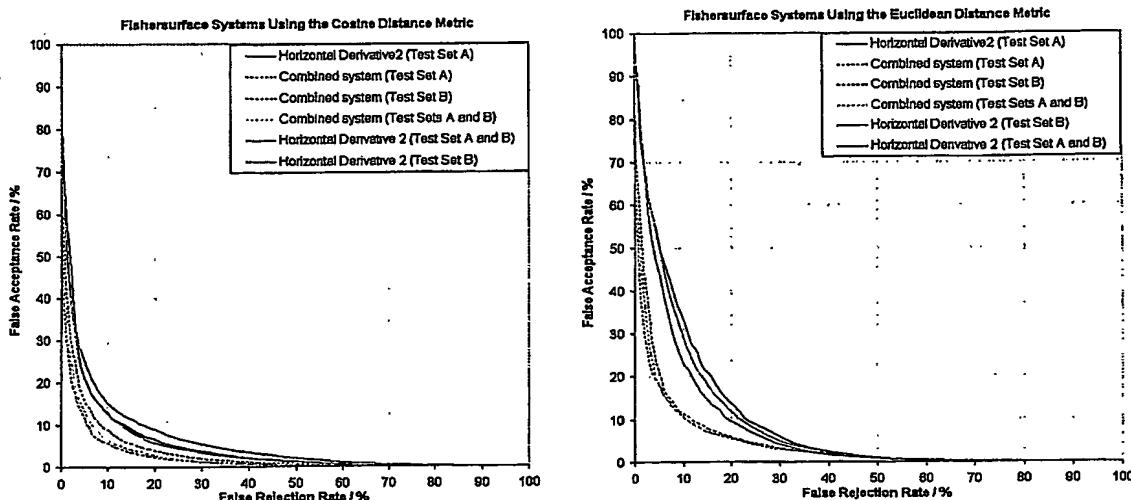
If *EER* has not increased:

Remove new *dimension* from *combined surface space*

Save *combined surface space* ready for evaluation

The figure below shows which dimensions from which surface space were selected using the above algorithm, for inclusion in two combined systems: one 10 using the euclidean distance metric and the other using the cosine distance metric.

We now compare the combined surface space systems with the optimum individual system, using both the cosine and euclidean distance measures.



The error curves above show the results obtained when the optimum single fishersurface system and combined fishersurface system are applied to test set A (used to construct the combined system), test set B (the unseen test set) and the full test set (all images from sets A and B) using the cosine and euclidean distance metrics. We see that the combined systems (dashed lines) do produce lower error rates than the single systems for both the cosine and Euclidean distance measure. The optimum system can be seen as the fishersurface combination using the cosine distance, producing an EER of 7.2% 9.3% and 8.2% for test set A, B and A and B respectively.

In one aspect, embodiments of the invention as described above apply the use of a fisherface method to the recognition of images, preferably natural images, preferably faces, and preferably human faces.

In another aspect, one or more pre-processing methods are applied prior to use of a fisherface method.

In another aspect, pre-processed data is combined prior to use of a fisherface method.

As indicated above, methods as disclosed herein may be combined with advantage with those disclosed in our prior, pending application GB0323662.7.

In this specification, the verb "comprise" has its normal dictionary meaning, to denote non-exclusive inclusion. That is, use of the word "comprise" (or any of its derivatives) to include one feature or more, does not exclude the possibility of also including further features.

The reader's attention is directed to all and any priority documents identified in connection with this application and to all and any papers and

documents which are filed concurrently with or previous to this specification in connection with this application and which are open to public inspection with this specification, and the contents of all such papers and documents are incorporated herein by reference.

5 All of the features disclosed in this specification, and/or all of the steps of any method or process so disclosed, may be combined in any combination, except combinations where at least some of such features and/or steps are mutually exclusive.

10 Each feature disclosed in this specification may be replaced by alternative features serving the same, equivalent or similar purpose, unless expressly stated otherwise. Thus, unless expressly stated otherwise, each feature disclosed is one example only of a generic series of equivalent or similar features.

15 The invention is not restricted to the details of the foregoing embodiment(s). The invention extends to any novel one, or any novel combination, of the features disclosed in this specification, or to any novel one, or any novel combination, of the steps of any method or process so disclosed.

References:

1. Heseltine, Pears, Austin.: Three-dimensional Face Recognition: An Eigensurface Approach. Under review for ICIP 2004.
- 20 2. Beumier, C., Achteroy, M.: Automatic 3D Face Authentication. Image and Vision Computing, Vol. 18, No. 4, (2000) 315-321
3. Gordon, G.: Face Recognition Based on Depth and Curvature Features. In Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Champaign, Illinois (1992) 108-110.
- 25 4. Chua, C., Han, F., Ho, T.: 3D Human Face Recognition Using Point Signature. Proc. Fourth IEEE International Conference on Automatic Face and Gesture Recognition, (2000) 233-8
5. Zhao, W., Chellaa, R.: 3D Model Enhanced Face Recognition. In Proc. Of the International Conference on Image Processing, Vancouver (2000)

6. Romdhani, S., Blanz, V., Vetter, T.: Face Identification by Fitting a 3D Morphable Model using Linear Shape and Texture Error Functions. The European Conference on Computer Vision (2002)
- 5 7. Blanz, V., Romdhani, S., Vetter, T.: Face Identification across Different Poses and Illuminations with a 3D Morphable Model. In Proc. of the 5th IEEE Conference on AFGR (2002)
8. Beumier, C., Acheroy, M.: Automatic Face Verification from 3D And Grey Level Clues. 11th Portuguese Conference on Pattern Recognition, 2000.
- 10 9. Heseltine, T., Pears, N., Austin, J.: Evaluation of image pre-processing techniques for eigenface-based face recognition. In Proc. of the 2nd International Conf. on Image and Graphics, SPIE Vol. 4875 (2002) 677-685
- 10 10. Heseltine, T., Pears, N., Austin, J.: Face Recognition: A Comparison of Appearance-based Approaches. In Proc. of the 7th International Conference on Digital Image Computing: Techniques and Applications, awaiting publication (2003).
- 15 11. A. Pentland, B. Moghaddam, T. Starner, "View-Based and Modular Eigenfaces for Face Recognition", *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, 1994.
12. Heseltine, T., Pears, N., Austin, J.: Combining multiple face recognition systems using Fisher's linear discriminant. In Proc. of the Defense and Security Symposium (2004).
- 20 13. Belhumeur, J. Hespanha, D. Kriegman, "Eigenfaces vs. Fisherfaces: Face Recognition using class specific linear projection", *Proc. of the European Conference on Computer Vision*, pp. 45-58, 1996.

01	02	03	04	05
				
Front neutral expression	Facing 45° up	Facing 45° down	Happy expression	Eyes closed
06	07	08	09	10
				
Angry expression	2 nd Neutral	Eyebrows raised	3 rd Neutral	Farther from camera

Fig. 1. Example face models taken from a 3D face database

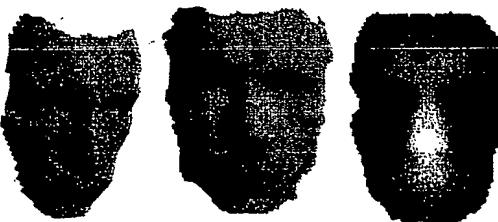


Fig. 2. Orientation of a raw 3D face model (*left*) to a frontal pose (*middle*) and facial surface depth map (*right*)



Fig. 3. Average depth map (*left most*) and first eight eigensurfaces

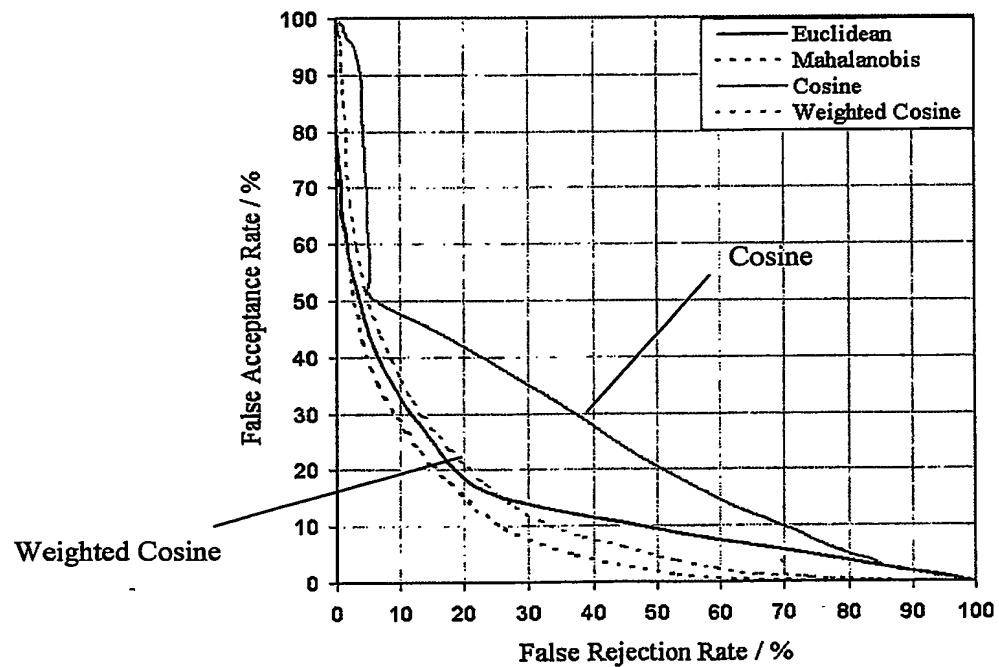


Fig. 4. Baseline 3D face recognition systems using facial surface depth maps and a range of distance metrics

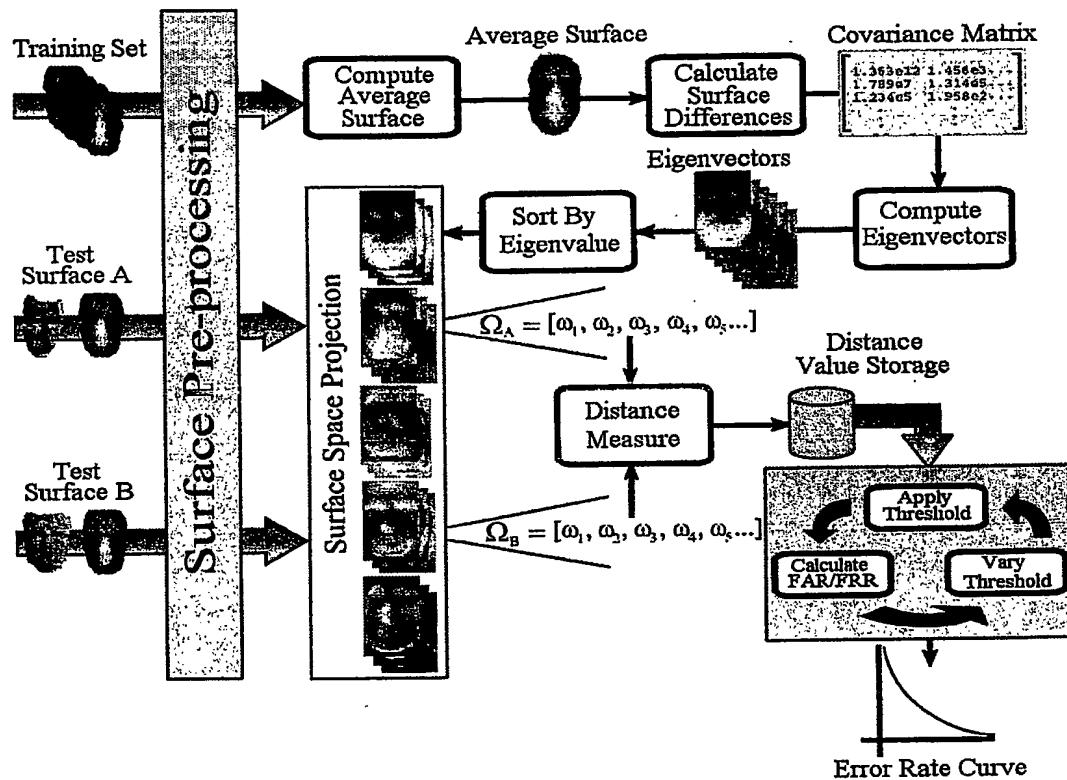


Fig. 5. Diagram of verification test procedure

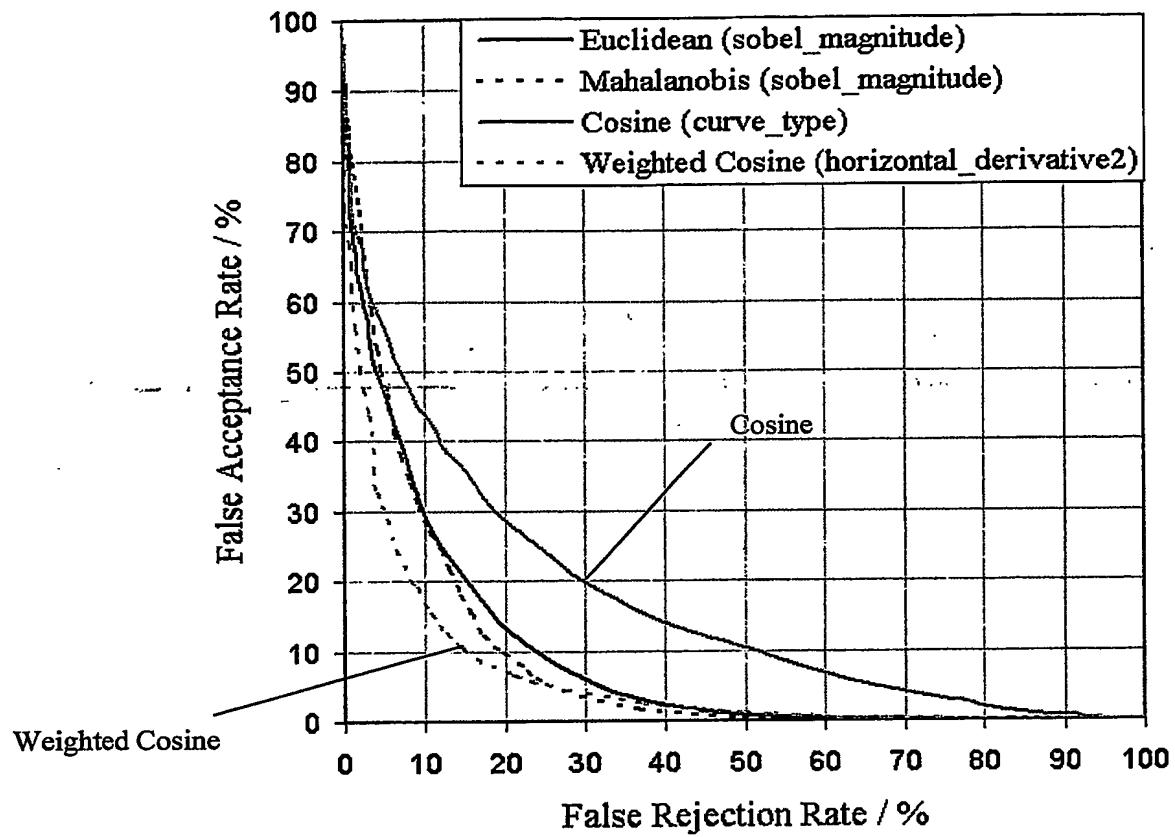


Fig. 6. Error rates of 3D face recognition systems using optimum surface representations and distance metrics.

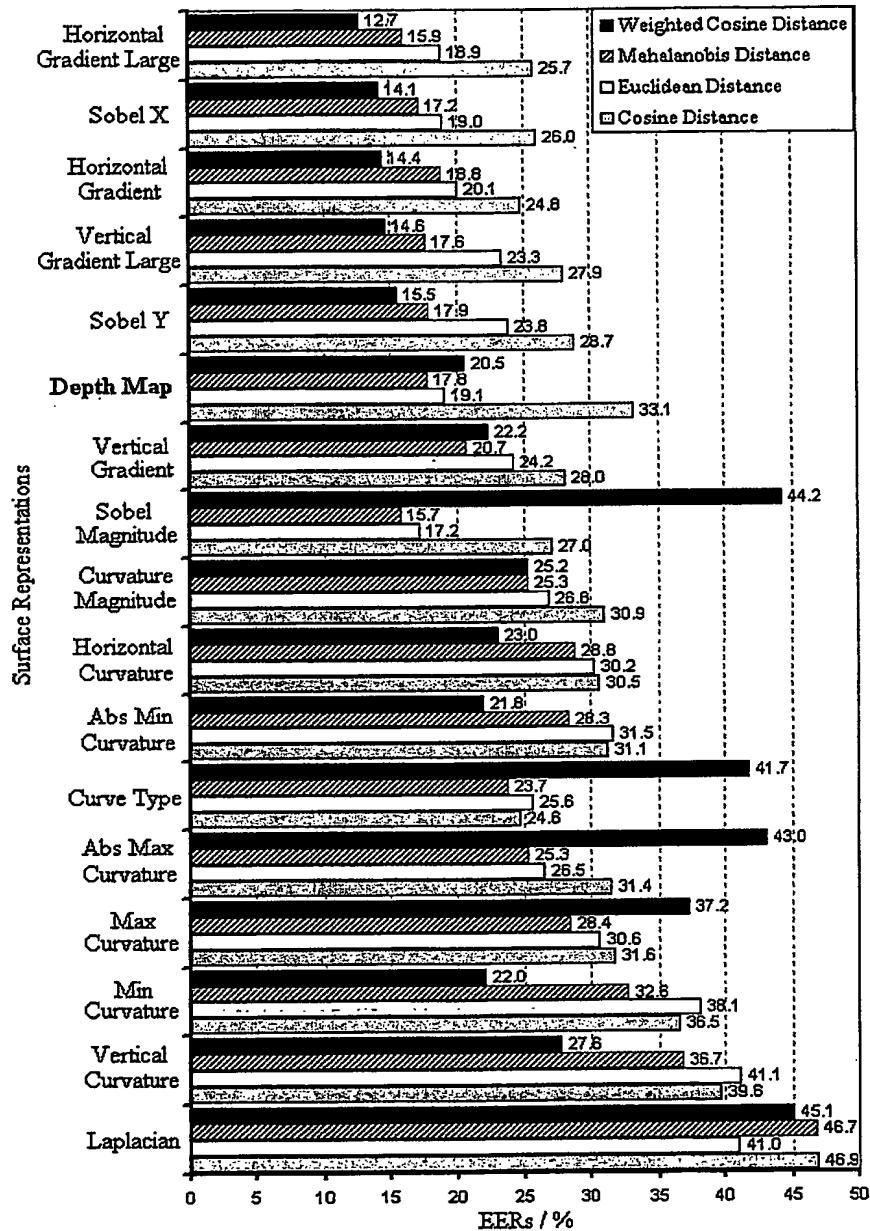
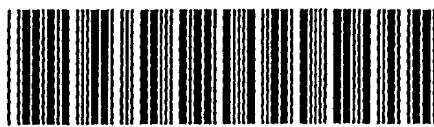


Fig. 7. Equal error rates of 3D face recognition systems using a variety of surface representations and distance metrics

Horizontal Gradient	Vertical Gradient	Horizontal Gradient Large	Vertical Gradient Large
	$-1 \ 1$		$-1 \ 1$
Applies the 2×1 kernel to compute the horizontal derivative	Applies the 1×2 kernel to compute the vertical derivative	Horizontal gradient calculated over a greater horizontal distance	Vertical gradient calculated over a greater vertical distance
Laplacian	Sobel X	Sobel Y	Sobel Magnitude
	$0 \ 1 \ 0$ $1 \ -4 \ 1$ $0 \ 1 \ 0$		$-1 \ 0 \ 1$ $-2 \ 0 \ 2$ $-1 \ 0 \ 1$
An isotropic measure of the second spatial derivative	Application of the sobel derivative filter in the horizontal direction	Application of the sobel derivative filter in the vertical direction	The magnitude of the X and Y sobel derivatives
Horizontal Curvature	Vertical Curvature	Curvature Magnitude	Curve Type
			
Applies the sobel X kernel twice to calculate the second horizontal derivative	Applies the sobel Y kernel twice to calculate the second vertical derivative	The magnitude of the vertical and horizontal curvatures	Segmentation of the surface into 8 discreet curvature types
Min Curvature	Max Curvature	Abs Min Curvature	Abs Max Curvature
			
The minimum of the horizontal and vertical curvature values	The maximum of the horizontal and vertical curvature values	The minimum of the absolute horizontal and vertical curvatures	The maximum of the absolute horizontal and vertical curvatures

Figure 8. Brief descriptions of surface representations with the convolution kernels used.

John
PCT/EP2004/052502



**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.